

Bayesian Methods for Bounding Single-Event Related Risk in Low-Cost Satellite Missions

R. L. Ladbury and M. J. Campola

Abstract—We develop single-event risk Prior probability distributions based on historical and heritage data. The Priors can be used to bound single-event effects risk for testing, part selection and design.

Index Terms—probabilistic risk assessment, radiation effects, reliability estimation, quality assurance, radiation hardness assurance.

I. INTRODUCTION

Cost constraints and microelectronic scaling have made low-cost platforms among the fastest growing sectors of the space industry. While such simple, cheap platforms facilitate rapid development and make satellite missions feasible for universities and small research laboratories, radiation effects threaten these efforts, and radiation hardness assurance (RHA) costs often strain budgets and schedule for such low-cost, risk-tolerant missions. This has proven especially true for threats arising from single-event effects (SEE).

Unfortunately, conventional SEE risk analysis methodologies have proven difficult to adapt to the budgets and risk tolerance of these missions. Such methodologies require that risks be bounded using empirical data, and test costs for obtaining such data are high. As such, many low-cost satellite projects have no comprehensive RHA strategy beyond “best effort.”

Absent rigorous RHA, projects lack appropriate risk metrics for component selection, design choices, derating and operational constraints. A risk mitigation method that bounds risk at all design stages could benefit low-cost, risk-tolerant missions by providing a metric for making such decisions as well as prioritizing decisions for allocating scarce testing resources to achieve maximum risk reduction. We previously developed Bayesian risk reduction strategies for Total Ionizing Dose (TID) degradation [1]. Here we develop similar Bayesian methods for SEE, which represent a more life-limiting threat for many low-cost, short missions in benign radiation environments. We first outline the strategy and then

apply it to specific radiation risks, including single-event latchup (SEL) in commercial analog-to-digital and digital-to-analog converters (ADCs and DACs) and single-event transients (SETs) from operational amplifiers (op amps).

II. BAYES’ THEOREM AND RISK

The goal of this RHA method is to bound SEE failure probabilities for a component in its application. Because such bounds can be thought of as subjective probabilities (that is, they depend on the state of our knowledge when we calculate them), this task is well suited to Bayesian probability. As with any SEE risk methodology, we would ideally base predictions of flight performance on the most specific and representative data—SEE test data for the flight-lot of the flight part type, or at least on a contemporaneous lot of that part. In the absence of such specific, controlled data, one could also look at in-flight performance of the same part type on previous missions—so-called heritage data. While heritage data provide constraints on the performance of the flight part types in which we are interested, such data also pose challenges, as discussed below. If such specific data were not available, the next level to look at would be SEE performance of similar parts fabricated in the same process. Unfortunately, since there is no way of telling whether the flight parts will perform better or worse than these similar parts, conventional RHA methodologies can only use such similarity data as a rough guide. Not only can the Bayesian approach outlined here place a quantitative bound on likely performance, it can do so for any desired confidence level (CL) and success probability P_s , allowing the RHA method to be tailored to the risk tolerance of the program.

For similarity data, we cannot speak directly to the performance of the flight parts. However, since the flight parts do belong to the larger class of all such parts fabricated in the same process, we can bound the performance of flight parts as long as they perform no worse than, for example, the median part, or the 90% worst case part in that class.

For heritage data, variability arises from differences in the environment and inherent Poisson fluctuations as well as in uncertainties about the charge collection geometries of flight parts.

Bayesian probability deals with changes in the probability that proposition x is true as we gather new data A . Prior to the new data, the probability that x is true is given by the Prior probability distribution, or “Prior” $P_0(x)$. Adding the new data

Manuscript received August 1, 2013. The authors thank the NASA GSFC IRAD program for support of this research and Sana Rezgui of LTC for tirelessly fielding questions about fabrication processes and design.

R. L. Ladbury is with the NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA (phone: 301-286-1030; fax: 301-286-4699; e-mail: raymond.l.ladbury@nasa.gov).

M. J. Campola is with the NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA, micael.j.campola@nasa.gov.

A gives us the posterior probability distribution $P(x|A)$ (read “ P of x given A ”)

$$P(x|A) = \frac{P(A|x)P_o(x)}{P(A)} = \frac{P(A|x)P_o(x)}{\int P(A|x')P_o(x')dx'} \quad (1)$$

where $P(A|x)$ is the likelihood of our data A given x is true, and the denominator is the probability of A integrated over all possible propositions x' , where x' includes x . An example of a proposition x would be that variability of SET rates in op amps fabricated in a given process is described by a lognormal distribution with lognormal mean and standard deviation m and s , respectively. If we have lots of prior experience with such parts, the Prior $P_o(x'=\{m,s\})$ may be peaked sharply around $x=\{m,s\}$, and it would take a lot of contradictory data to make the Posterior deviate significantly from the Prior. Alternatively, if we have little experience, the Prior may be broad and flat. In this case, the Posterior will look like the likelihood $P(A|x)$ and be determined almost exclusively by the data. Initial Priors may be based on any data pertinent to and representative of the flight parts we are considering—including similarity or heritage data or even expert opinion. Once a posterior distribution is calculated, it serves as the Prior for the next time data are added.

III. DATA SOURCES

We summarize the data used in this study in Tables I and II. We have drawn the data for the following analyses from publicly available sources, including Radiation Effects Data Workshop papers [2, 6] and publicly accessible [3, 4, 5] websites. We determined the process in which the parts were fabricated by consulting the vendors’ websites or by sending inquiries to the vendors. As an illustration of the method, we treat two SEE risks: SETs in op amps and SEL in ADCs and DACs. Table I shows the part number, manufacturer (Linear Technology Corp. (LTC) or Analog Devices Inc. (ADI)) and process, SET rate at geostationary orbit due to galactic cosmic rays (GCR), worst-case (WC) transient duration and amplitude and the application conditions for the test. Table II shows SEL onset linear energy transfer (LET₀) and limiting cross sections (σ_{lim}) for ADI ADCs and DACs fabricated in their 0.6 μm complementary metal oxide semiconductor (CMOS) process taken at Brookhaven National Laboratory (BNL) or Texas A&M University’s cyclotron facility (TAMU).

The main challenges we face with our current method are gathering sufficient data to ensure adequate representation of similar parts across a particular fabrication process and developing metrics that are sufficiently general to apply across the entire family of devices but still yield meaningful bounds on likely performance of flight parts. For instance, while SET rates and widths in op amps yield sufficiently compact distributions for meaningful modeling, transient amplitudes depend on application conditions such as supply voltage. As such, it is preferable to normalize worst-case transient amplitude to supply voltage, yielding a ratio defined from 0 to 1 that can be modeled with a Beta distribution.

TABLE I: SET DATA FOR ADI AND LTC OP AMPS

Part Number	Vendor/Process	SET Rate (day ⁻¹)	WC Duration (μs)	WC Amplitude (V)	Supply Voltage -- Output Voltage (V)
RH108 [2]	LTC - RH Bipolar	0.011	7	4.5	0/12 -- 6
RH118 [2]	LTC - RH Bipolar	0.046	> 5	2	-8/+8 -- 0
RH1014 [2]	LTC - RH Bipolar	0.029	> 35	Rail-to-Rail	-12/+12 -- 6
RH1499 [3]	LTC - RH Bipolar	0.0082	13	4.4	-15/+15 -- 4
RH1078 [4]	LTC - RH Bipolar	0.04	30	2	0/10 -- 2
OP27 [5]	ADI - Bipolar > 2.5 μm	0.036	10	10	-10/+10 -- Varied
OP113 [2]	ADI - Bipolar > 2.5 μm	0.01	> 2	1.7	0/5 -- 2.55
OP270 [2]	ADI - Bipolar > 2.5 μm	0.35	1.2	3	-5/+5 -- 0.42
OP400 [2]	ADI - Bipolar > 2.5 μm	0.56	10	3	-12/+12 -- ± 2.4
OP05 [5]	ADI - Bipolar > 2.5 μm	--	12	--	--
OP15 [5]	ADI - Bipolar > 2.5 μm	--	15	--	--

TABLE II: SEL DATA FOR ADI 0.6 μm CMOS ADCs AND DACs

Part #	Process	Onset LET (MeVcm ² /mg)	Test Facility	σ_{lim} (cm ²)
AD5305	0.6 μm	LET<11	BNL	??
AD7472	0.6 μm	11<LET<15	BNL	5.00E-04
AD7476	0.6 μm	LET<60	BNL	??
AD7664	0.6 μm	6	BNL	2.50E-04
AD7714	0.6 μm	15<LET<24	BNL	6.00E-04
AD9260	0.6 μm	8	BNL	??
AD5334	0.6 μm	5	TAMU	5.00E-04
AD7664	0.6 μm	4	TAMU	1.00E-03
AD7675	0.6 μm	4	TAMU	3.00E-05
AD7714	0.6 μm	20<LET<22	TAMU	1.00E-04

IV. BOUNDING SET RATES FOR OP AMPS

Since SET rates are positive definite, we assumed that variability of SET rates across op amps in a particular process followed a lognormal distribution. We started with a broad, uninformative Prior over all possible combinations of lognormal mean and standard deviation. We then calculated the likelihood of the SET rate data for each vendor in Table I for each parametric combination and updated our Prior with this likelihood. This resulted in the posterior probability distributions shown in Fig. 1 for the ADI OP series and in Fig. 2 for the LTC RH series op amps. The probability for a parametric pair is the probability that SET rates across op amps in the process follow a lognormal distribution with those parameters. The lognormal mean m equals the logarithm of the median SET rate, and the lognormal standard deviation s is related to the ratio of the arithmetic standard deviation σ and mean μ :

$$s = \sqrt{\ln(1 + \sigma^2 / \mu^2)} \quad (2)$$

The parametric combination with the highest probability gives the best fit to the data, so the expected SET rate for “typical” op amps in the process is the mean rate for the most probable fit parameters. For ADI parts, this is 0.055 SETs per day. The rate at 90% cumulative probability for this distribution corresponds to the expected rate for the 90% worst-case (WC) op amp in the process, in this case 0.089 SETs per day.

Moving away from the most probable parametric combination, we can draw contours containing increasing

cumulative probabilities. These cumulative probabilities correspond to increasing confidence that the actual best fit for the parts in the fabrication process are contained within the said contour. By considering distributions within a particular probability contour, we can bound the rate more conservatively, taking into account possible biases introduced by the sample of parts for which we have data. For example, if we take the maximum 90% WC SET rate (corresponding to 90% cumulative probability of the lognormal) yielded by any distribution within the 90% cumulative probability contour in Fig. 1, this will, with 90% confidence, bound the SET rate for 90% of op amps in the process—0.19 SETs per day for ADI OP series amplifiers. Similarly, using the posterior distribution in Fig. 2, we found the best estimate for the SET rate in a typical LTC RH series op amp was 0.027 SETs per day. The 90% WC op amp (again using the most probable distribution given our data) was 0.05 SETs per day. With 90% confidence, less than 10% of LTC RH series op amps would experience SETs due to heavy ions at a rate of 0.13 SETs per day.

A significant advantage of this method is that it allows the confidence level and probability of success to be selected according to the risk tolerance of the mission and the application. Moreover, by making these levels of assurance concrete, it facilitates comparison of radiation risk with those arising from other sources

In the analysis for both of the above datasets, we have chosen to model the distribution of transients as lognormal for purposes of illustration. However, one can choose other appropriate distribution forms, such as Weibull, to assess the distribution dependence of results. Generally, as long as one starts with an equally broad Prior, the distribution form that yields the highest peak probability will yield the best fit to the data.




Fig. 1: Prior probability distribution of SET rates expected for op amps in ADI's OP series modeled as a lognormal distribution. The most probable lognormal mean and standard deviation are -3 and 0.45.




Fig. 2: Prior probability distribution of SET rates expected for LTC RH series op amps modeled as a lognormal distribution. The most probable lognormal mean and standard deviation are -3.8 and 0.625.

V. BOUNDING SET DURATION AND AMPLITUDE

Similarly, we estimate worst-case transient duration and amplitude across op amps in a process. Assuming transient durations vary lognormally across part types for both ADI and LTC op amps, the resulting Prior distributions are shown in Fig. 3 for ADI and Fig. 4 for LTC. Table III shows the most probable and 90% CL bounds on transient width for typical and 90% WC op amps for both families of op amps.

TABLE III: BOUNDING SET WIDTHS FOR ADI AND LTC OP AMPS

Process	Typical	90% WC	Typical@90% CL	90%WC@90%CL
ADI	15 μ s	<35 μ s	143 μ s	<218 μ s
LTC	22 μ s	<35 μ s	75 μ s	<170 μ s

Despite the limited data in Table I, the best-fit typical and 90% WC transient durations in Table III are consistent with those seen from op amps under worst-case conditions. When we consider behavior possible at the 90% confidence level, the limited data are evident in the much longer transient durations. This illustrates that the approach taken here is conservative. As long as data are representative of flight parts, the method supplies a conservative bound even if data are limited. This approach encourages the analyst to gather as much data as possible to narrow the Prior, ensures efficient but not overly optimistic use of the information and allows the required probability of success and confidence level to be tailored to mission risk tolerance.




Fig. 3: Prior probability distribution for SET width for ADI's OP series op amps modeled as a lognormal distribution.




Fig. 4: Prior probability distribution for SET width for LTC's RH series op amps modeled as a lognormal distribution.

Modeling worst-case transient amplitudes is a more difficult proposition, given the wide range of operating conditions (*e.g.*, supply and input voltages, feedback, *etc.*) for these parts. One approach that renders the data more tractable is to model the ratio of the maximum transient amplitude to the supply voltage. This yields a number between 0 and 1, which can be modeled as a Beta distribution. However, because rail-to-rail transients are common in op amps, even the most probable fit to the data yields a significant probability of transient amplitude at or near the rail. For LTC RH series op amps, the 90% worst-case part yields transients $\sim 90\%$ of the rail. For ADI parts, the 90% worst-case part yields even larger transients—95% of the rail. These estimates would only be at the 50% confidence level. This suggests that assuming the transient amplitude can reach the supply voltage is not only prudent, but realistic.

VI. BOUNDING SEL RISK

Destructive SEE, such as SEL, are among the most serious risks facing low-cost, short missions. Unfortunately, bounding risk for SEL is complicated. Limiting cross sections in Table II vary over a factor of 30, making it impossible to form a tight distribution for SEL rates. SEL susceptibility depends not just




Fig. 5 Prior probability distribution of SEL onset LETs for ADCs and DACs fabricated in ADI's 0.6- μm CMOS process. The most probable lognormal mean and standard deviation are 1.45 and 0.11, respectively, yielding an expected $\text{LET}_0 = 4.3 \text{ MeV}\cdot\text{cm}^2/\text{mg}$.

on process characteristics, but also on how aggressive the design is for the process. We have examined data for several CMOS generations for both ADI and LTC parts, and where sufficient data exist to reveal it, even LET_0 seems to be distributed bimodally. For the ADI 0.6 μm process presented as an example here, one cluster is in the single digits and another is between 11 and 22 $\text{MeV}\cdot\text{cm}^2/\text{mg}$. SEL susceptibility also depends on ion range as well as LET, so short-range ions such as those at BNL serve only as upper bounds for LET_0 and lower bounds for limiting cross section.

Since SEL rates are not amenable to modeling due to their broad spread, we instead look at variability of LET_0 for SEL. Since we do not know which mode the flight parts would fall into, our bound must assume they would belong to the lower mode. Considering only those parts belonging to the lower mode from Table II, we update a broad uninformative Prior over possible lognormal parameters using a likelihood where TAMU onset LETs are treated as point estimates and BNL data are treated as upper limits on LET_0 . Fig. 5 shows the resulting Prior for variability in onset LET modeled as a lognormal. Proceeding as before for SETs, we consider the most probable lognormal distribution given our data. The most likely typical onset LET for this process is $4.3 \text{ MeV}\cdot\text{cm}^2/\text{mg}$, with the 90% WC part having $\text{LET}_0 > 3.7 \text{ MeV}\cdot\text{cm}^2/\text{mg}$. With 90% confidence, the corresponding onsets would be > 3.9 and $> 2.7 \text{ MeV}\cdot\text{cm}^2/\text{mg}$.

Once we have a bound for LET_0 at the desired confidence level and probability, we must use it to bound the SEL rate. One way to do this is to use the Figure of Merit (FOM) [7] approach to estimate the SEL rate R :

$$R = C_E \times \frac{\sigma_{\text{lim}}}{\text{LET}_{0.25}^2} \quad (3)$$

where σ_{lim} is the limiting cross section, $\text{LET}_{0.25}$ is the LET where the cross section reaches 25% of σ_{lim} and C_E is a constant determined by the mission radiation environment. For the purposes of bounding the rate, we can substitute LET_0 for $\text{LET}_{0.25}$, since we know this will over-estimate the rate.

For σ_{lim} , we can take the maximum saturated cross section from Table II—about 0.001 cm^2 . For the International Space Station (ISS) orbit, this corresponds to a rate of less than about once in 7.5 months with 90% confidence, while in a geostationary orbit (GEO) the bound would be more than once per month. Thus, while an arbitrary ADC or DAC might be acceptable for a short mission on the ISS, it would likely be problematic in harsher environments.

In addition to the $0.6 \mu\text{m}$ CMOS process, we have also examined the $0.8\text{-}2 \mu\text{m}$ CMOS process in a similar manner and found with 90% confidence that $> 90\%$ of ADCs and DACs would have $\text{LET}_0 > 3.5 \text{ MeV}\cdot\text{cm}^2/\text{mg}$. For ADCs in LTC's 120-nm CMOS process, with 90% confidence, $> 90\%$ would have $\text{SEL LET}_0 > 3.7 \text{ MeV}\cdot\text{cm}^2/\text{mg}$. SEL response in both of these processes appears to be bimodal. Perhaps most notable, of the 5 parts in ADI's BiCMOS process for which we could find data, none exhibited susceptibility to SEL.

VII. UPDATING PRIORS WITH HERITAGE

For the above method to be truly Bayesian, new data must be generated regularly to update the Priors. If a mission cannot afford to test critical parts, the mission success or failure of those parts becomes the data with which we must update our Prior. If the parts succeed, we can use the lack of observed failures to constrain the failure rate. Since SEE are Poisson processes, an observation of 0 failures during the mission implies the mean or expected number of failures is less than 2.3 with 90% confidence. Thus, for an identical mission in an identical environment, the failure rate is less than $2.3/T_m$ (90% CL), where T_m is the mission duration.

If the heritage and new missions are in different environments, we must find a way to equate a heritage mission duration to an equivalent duration for our desired reference mission environment. One possible strategy is to use Eq. 3 to estimate appropriate weights. If C_{EH} and C_{ER} are the FOM coefficients for the Heritage and reference missions, then each year of the heritage mission counts $W_{HR} = C_{EH}/C_{ER}$ for the reference mission environment. By this method, each year on the ISS counts about 1 month in a geostationary (GEO) or planetary environment. However, we note that the FOM coefficients represent an overall best fit over many different parts with different charge collection geometries, and the exact weight will depend on the geometries of charge collection volumes as well as the specific environments. To illustrate this effect, we have estimated rates over a broad range of device geometries for GEO, a polar orbit (750 km, 98 degrees inclination, sun-synchronous) and the space station orbit (500 km, 51.6 degrees inclination). We then took the ratio of the geostationary rate to the polar and space station rates for each device geometry. Fig. 6 illustrates the range of these ratios, along with how they compare to the ratio of the FOM environmental coefficients for the same orbits.

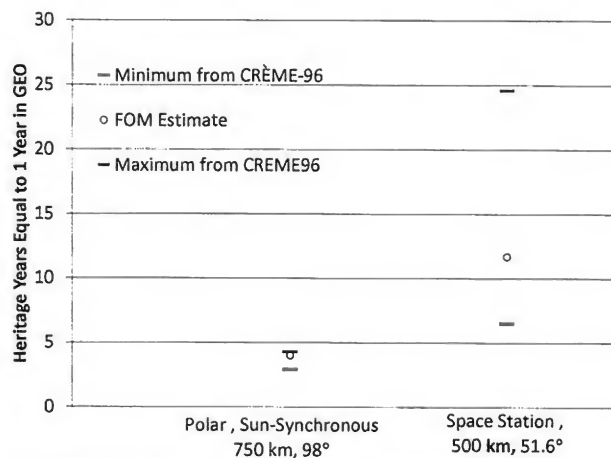


Fig. 6: Extrapolating heritage mission SEE performance to missions in other radiation environments requires estimating weights that equate time spent in the heritage environment with time spent in the new mission environment. The weights also depend on the charge collection geometries of the parts. The importance of this effect differs for different heritage and reference missions—by less than 40% between a polar heritage mission and a GEO mission and by a factor of almost 4 in going from a space station environment to a GEO mission.

For the polar orbit, the range of weights is tight (2.9 to 4.3), and the FOM coefficient ratio provides a good estimate of the weight for all device geometries. However, for the space station environment, the CREME96 weights span a factor of $\sim 4\times$ (6.5 to 24.6), with the FOM ratio lying roughly in the middle (~ 11.7) [8]. This exercise illustrates several important considerations when using heritage data. First, the more dissimilar the environments (especially as regards geomagnetic shielding), the more likely device geometry is to play an important role in determining the proper weight for the heritage mission. Second, if we have information about possible device geometries—e.g., from similarity data—we may be able to restrict the range of weights that are likely for our flight parts. This allows more confident use of heritage data, while also ensuring that it is not given too much weight. Third, it is important to avoid giving too much weight to heritage data, especially for benign environments such as the ISS. Considering Poisson fluctuations and the range of weights for conversion of ISS years to GEO years, all we could conclude as a result of a 5-year mission on board the ISS is that with 90% confidence we would expect the system to last longer than 1.1 months in GEO.

VIII. CONCLUSIONS

For short-duration, low-cost missions, SEE are often the dominant radiation risk. Unfortunately, the lack of a SEE risk analysis method that allows risk to be tailored to the tolerance of such missions makes it difficult to bound radiation risk, complicating part selection, circuit design and mission planning. This can result in overly conservative designs that limit performance or in premature failure. What is more, in many cases, the causes of failure are never determined, complicating the task of improving subsequent missions. The preceding analysis seeks to rectify this shortcoming by

introducing a procedure for deriving Bayesian Prior probability distributions for SEE from similarity and heritage data. Provided the data on which the Prior is based are representative of flight parts, the Prior can be used to bound flight part SEE risk. The analysis also illustrates some of the challenges facing such a method, chief among which are finding sufficient representative data for a representative Prior, assessing the impact of application specific testing and developing metrics and models that yield compact distributions and meaningful predictions for the mission. For some datasets (e.g., those used for SET rates and durations here), modeling is straightforward. For others (e.g., for transient amplitudes), the analysis reveals that the worst case (e.g., rail-to-rail transients) is, in fact, a realistic possibility and should be planned for. As we saw for the consideration of SEL risk, parts of the risk may be modeled (LET_0 variability within a process), but additional assumptions are required to fully bound risk.

The use of heritage data in qualifying parts for new missions requires special consideration, as the proper weighting of such data depends not just on the heritage and new mission environments, but also on the device's charge collection geometry. Naïvely extrapolating from one environment to another could result in underestimation of the SEE rate by a factor of 2 or more.

Ultimately, for the approach outlined here to be truly Bayesian, we must have new data to update our Prior. If SEE testing is not an option, the success or failure of the parts during the mission will serve as our data. This means we must know not just whether the mission succeeded or failed, but if it failed, whether the parts in which we are interested caused the failure. This minimum standard ensures that the method outlined here is self-correcting and improves continually over time as well as provides assurance of improved platform performance over time.

Although the methods outlined here are designed to provide sufficient flexibility to capitalize on the increased risk tolerance of low-cost, small satellite platforms, the techniques should also be useful for conventional satellites. The Priors allow risk to be bounded at any level of confidence desired by the project even prior to the availability of SEE test data for flight parts. Such bounds allow test efforts to be prioritized to maximize risk reduction. Moreover, by considering test results in the context of similar parts in the same process, the method makes it possible to flag outliers, providing early alerts to possible counterfeit parts, process changes or other out-of-family results that deserve added scrutiny. The method also provides a way of using heritage data rationally that avoids over-optimistic assumptions about the value of such flight experience. Finally, the Bayesian methods can easily be extended to allow inclusion of results from physical models or trends across different processes or even technological node.

REFERENCES

- [1] Ladbury, R. and Triggs, B., "A Bayesian Approach for Total Ionizing Dose Hardness Assurance," *IEEE Trans. Nucl. Sci.*, Vol. 58, No. 6, pp. 3004 – 3010, 2011.
- [2] George, J. *et al.*, "Single event transients in operational amplifiers," *IEEE Radiation Effects Data Workshop*, 2005, pp. 8–12, 2005.
- [3] Buchner, S and Forney, J. "Single Event Transient Testing of RH1499," http://radhome.gsfc.nasa.gov/radhome/papers/L061806_RH1499.pdf, 2006.
- [4] Guerre, F. X "Heavy Ion Testing of the RH1078 Dual, Single Supply Precision Op Amp from Linear Technology," <https://escies.org/download/webDocumentFile?id=981>, 10/29/1998.
- [5] Poivey, C., "Application Notes for Analog Linear Devices," https://nepp.nasa.gov/files/11786/SET_Application_Draft.pdf, 2005.
- [6] Irom, F., Agarwal, S. G. "Compendium of Single-Event Latchup and Total Ionizing Dose Test Results of Commercial Analog to Digital Converters," *IEEE Radiation Effects Data Workshop* 2012, 14 pages, 2012.
- [7] Petersen, E. L., "The SEU figure of merit and proton upset rate calculations," *IEEE Trans. Nucl. Sci.*, Vol. 45, No. 6, pp.2550-2562, 1998.
- [8] Tylka, A. J. *et al.*, "CREME96: A Revision of the Cosmic Ray Effects on Micro-Electronics Code," *IEEE Trans. Nucl. Sci.*, Vol. 44, No. 6, pp.2150-2160, 1997.